



EEHPC Session 5: Power API HPC Controls and Monitoring

James H. Laros III

Sandia National Laboratories

<http://powerapi.sandia.gov>



Sandia
National
Laboratories

*Exceptional
service
in the
national
interest*





U.S. DEPARTMENT OF
ENERGY



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND2015-10111 C

What Did We Learn? (from our early research)





It's a
Fan!

It's
a Wall!

It's
a
Rope!


It's a
Spear!

It's
a Snake!

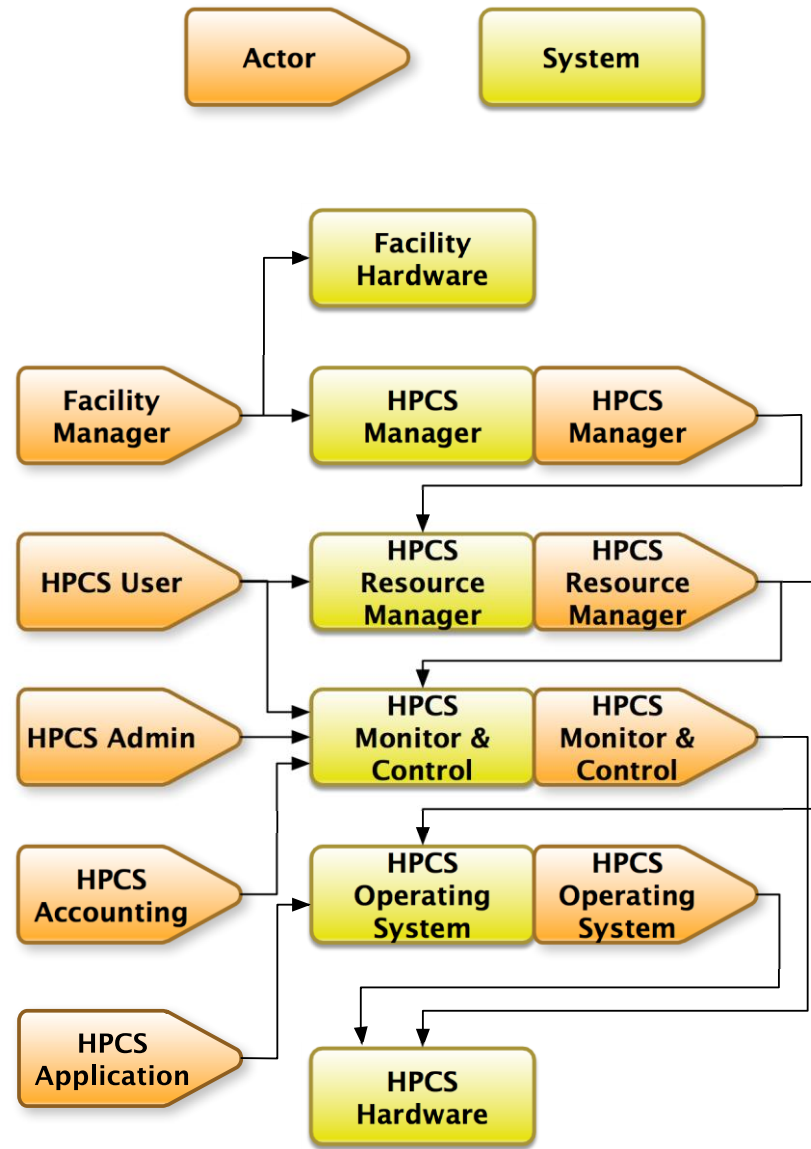
It's a
Tree!

Preliminary Investigation

- Sandia and National Renewable Energy Laboratory (NREL)
 - Sandia – Platform focused
 - NREL – Facility focused
- Primary Objectives
 - Identify SCOPE of Power API
 - Identify Roles
 - Identify Interfaces
- Product:
 - “Power/Energy Use Cases for High Performance Computing”
 - Use Case Approach
 - **Actor** (Roles) -> **System** pairs
 - Interaction between **Actor** and **System** defines an Interface
 - Facility to Component (hardware) **Scope**



Result



Version 1.1a Delivered

- Community needed a portable API for **measuring** and **controlling** power and energy
- Sandia developed Power API specification to fill this gap
- Provides portable power **measurement** and **control** interfaces, covers full spectrum of facility to component
- First production implementation in progress on Trinity (ATS1)
- Continued (increasing) community involvement and influence

SANDIA REPORT

SAND2015-6778
Unlimited Release
Printed September 2015

High Performance Computing - Power Application Programming Interface Specification Version 1.1a

James H. Laros III, David DeBonis, Ryan Grant, Suzanne M. Kelly, Michael Levenhagen, Stephen Olivier, Kevin Pedretti

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

<http://powerapi.sandia.gov>

System Description

PWR_ObjType

```
typedef enum {  
    PWR_OBJ_PLATFORM,  
    PWR_OBJ_CABINET,  
    PWR_OBJ_CHASSIS,  
    PWR_OBJ_BOARD,  
    PWR_OBJ_NODE,  
    PWR_OBJ_SOCKET,  
    PWR_OBJ_CORE,  
    PWR_OBJ_POWER_PLANE,  
    PWR_OBJ_MEM,  
    PWR_OBJ_NIC,  
    PWR_OBJ_INVALID  
} PWR_ObjType;
```

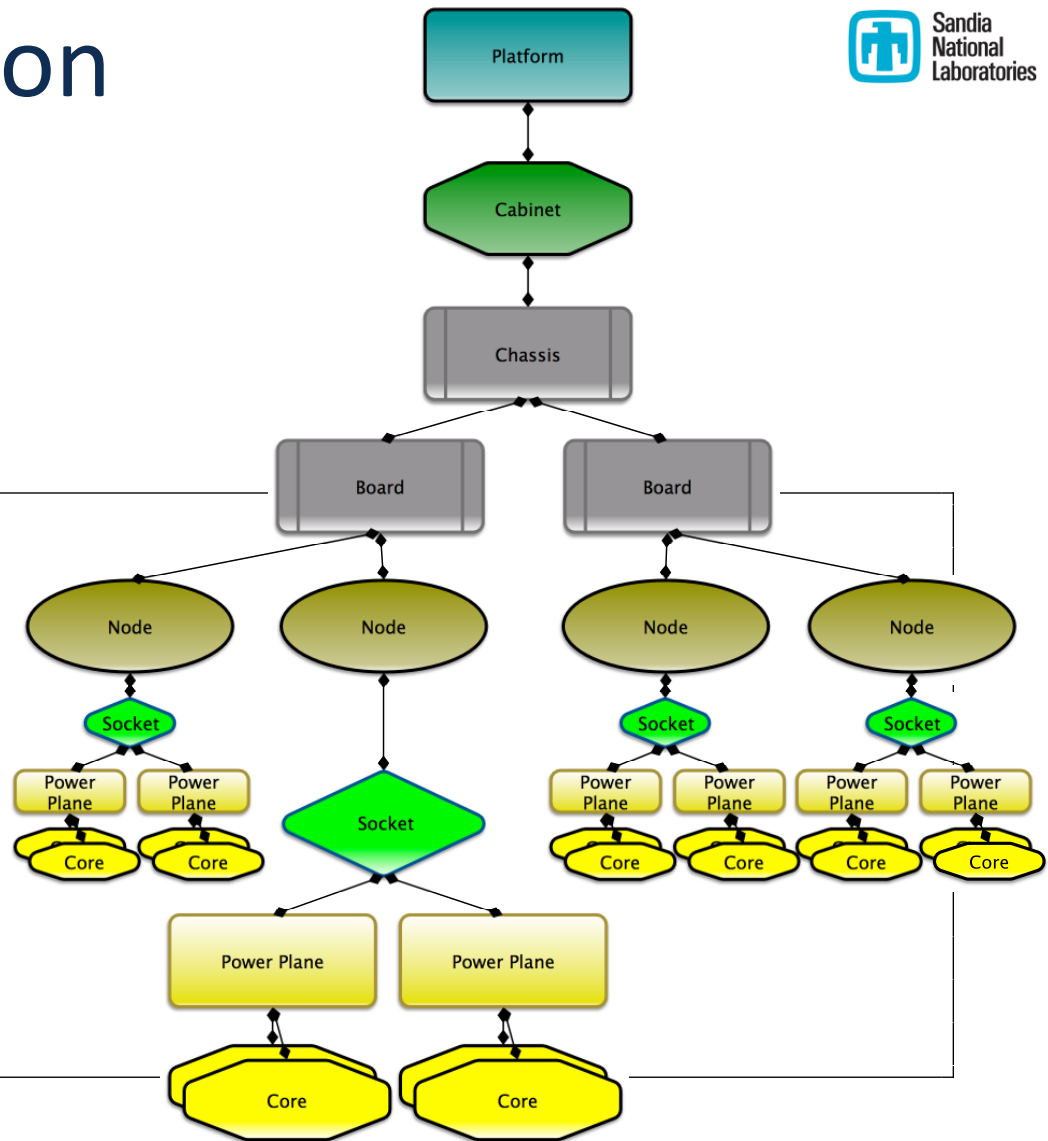
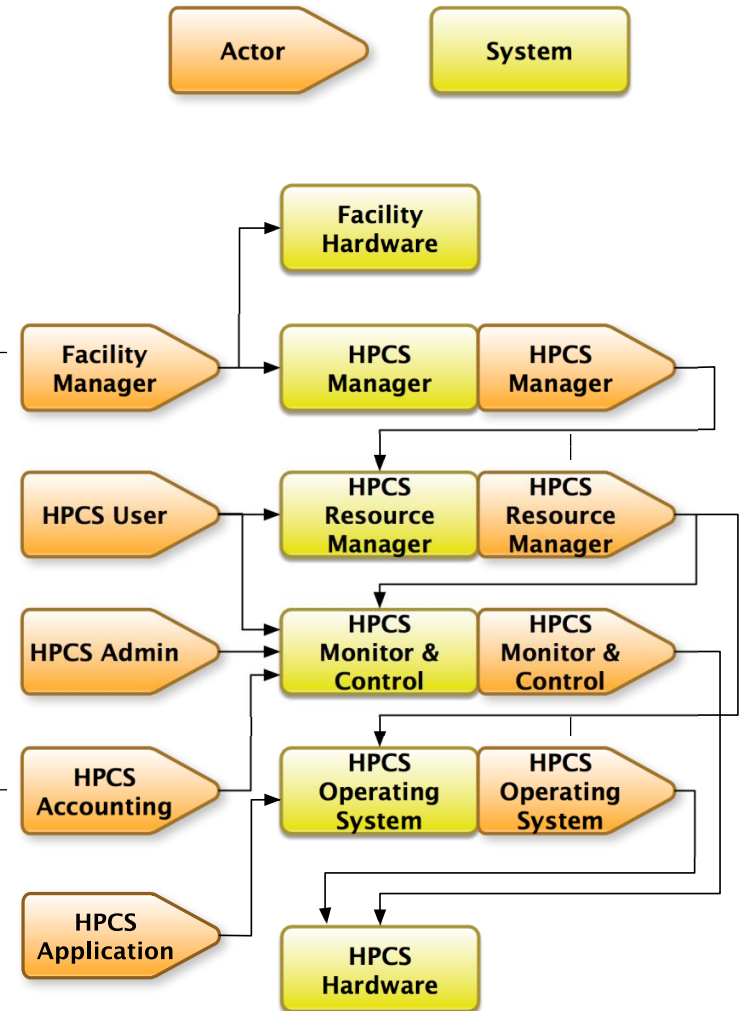


Figure 2.1: Hierarchical Depiction of System Objects

Roles

PWR_Role

```
typedef enum {  
    PWR_ROLE_APP, /* Application */  
    PWR_ROLE_MC, /* Monitor and Control */  
    PWR_ROLE_OS, /* Operating System */  
    PWR_ROLE_USER, /* User */  
    PWR_ROLE_RM, /* Resource Manager */  
    PWR_ROLE_ADMIN, /* Administrator */  
    PWR_ROLE_MGR, /* HPCS Manager */  
    PWR_ROLE_ACC /* Accounting */  
} PWR_Role;
```



Foundation: Measurement and Control

PWR_AttrName

```
typedef enum {  
    PWR_ATTR_PSTATE = 0, /* uint64_t */  
    PWR_ATTR_CSTATE, /* uint64_t */  
    PWR_ATTR_CSTATE_LIMIT, /* uint64_t */  
    PWR_ATTR_SSTATE, /* uint64_t */  
    PWR_ATTR_POWER, /* double, Watts */  
    PWR_ATTR_CURRENT, /* double, Amps */  
    PWR_ATTR_VOLTAGE, /* double, Voltage */  
    PWR_ATTR_MAX_POWER, /* double, Watts */  
    PWR_ATTR_MIN_POWER, /* double, Watts */  
    PWR_ATTR_FREQ, /* double, Hz */  
    PWR_ATTR_ENERGY, /* double, Joules */  
    PWR_ATTR_TEMP, /* double, Celsius */  
    PWR_ATTR_OS_ID, /* uint64_t */  
    PWR_ATTR_NUM_ATTRS,  
    PWR_ATTR_INVALID = PWR_ATTR_NUM_ATTRS,  
} PWR_AttrName;
```

Attribute Interface

MEASURE

```
int PWR_ObjAttrGetValue( PWR_Obj object,  
                        PWR_AttrName attr,  
                        void* buf,  
                        PWR_Time* ts);
```

CONTROL

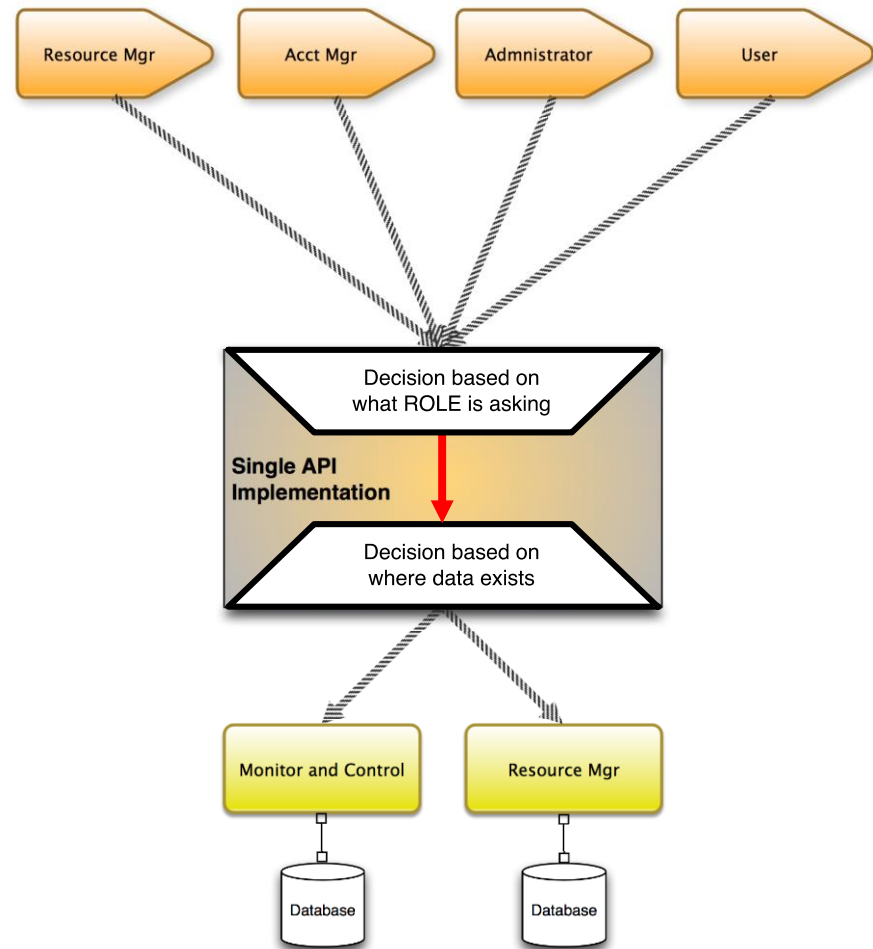
```
int PWR_ObjAttrSetValue( PWR_Obj object,  
                        PWR_AttrName attr,  
                        void* buf );
```

Other Core Features

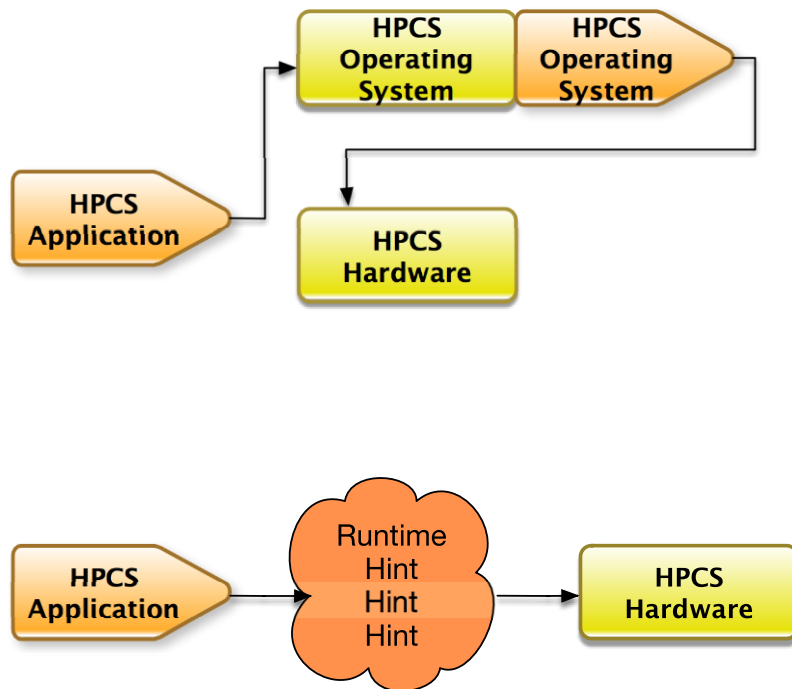
- Discovery and Navigation
 - Of system description
- Groups
 - Adds lots of flexibility and convenience
- Statistics
 - Real-time, historic and extensible
- Metadata
- Languages
 - V 1.0 – 1.1a Specified in C
 - V 1.x Python

Higher Level Interfaces

- Administrator:
 - Hard Power Capping
 - Bounds more granular
 - Power Aware Scheduling
- Resource Manager:
 - Power Aware Scheduling
 - HPC Tetris
- Accounting Manager:
 - Power/Energy Application Profiles
 - Feed into Power Aware scheduling
- User:
 - Understanding application power and energy characteristics
 - Application Hints



Higher Level Interfaces (cont.)



Application provides “hints”

- I’m going into a serial portion
 - Turn off cores
 - Accelerate active core
- Could be an intelligent dynamic run-time
- Research area

14

Questions?

- <http://powerapi.sandia.gov>
- Register on the reflector
- Get the current version of the spec
- Get the prototype source
- Other information as it develops
- Please get involved and help us improve the specification
- TEAM:

James Laros, Suzanne Kelly, Kevin Pedretti, Michael Levenhagen,
Ryan Grant, Stephen Olivier, David DeBonis